

A 3D fast algorithm for computing Lagrangian coherent structures via ridge tracking

Doug Lipinski^{a,c}, Kamran Mohseni^{a,b,c}

^a*Dept. of Mechanical and Aerospace Engineering, University of Florida, P.O. Box 116250, 231 MAE-A, Gainesville, FL 32611*

^b*Dept. of Electrical and Computer Engineering, University of Florida, P.O. Box 116200, 216 Larsen Hall Gainesville, FL 32611*

^c*This work was begun while the authors were affiliated with the Departments of Applied Mathematics and Aerospace Engineering Sciences at the University of Colorado - Boulder.*

Abstract

Lagrangian coherent structures (LCS) in fluid flows appear as co-dimension one ridges of the finite time Lyapunov exponent (FTLE) field. In three-dimensions this means two-dimensional ridges. A fast algorithm is presented here to locate and extract such ridge surfaces while avoiding unnecessary computations away from the LCS. This algorithm reduces the order of the computational complexity from $\mathcal{O}(1/dx^3)$ to about $\mathcal{O}(1/dx^2)$ by eliminating computations over most of the three dimensional domain and computing the FTLE only near the two-dimensional ridge surfaces. The algorithm is grid based and proofs of error bounds for ridge locations are included. The algorithm performance and error bounds are verified in several examples. The algorithm offers significant advantages in computational cost as well as later data analysis.

Keywords: Lagrangian coherent structures, fast algorithm, ridge tracking

1. Introduction

Lagrangian coherent structures (LCS) have seen increasingly popular use for visualizing and quantifying fluid structures and transport and mixing behavior in fluid flows. LCS were proposed by Haller and Yuan as the locally most repelling or attracting material lines in a flow [1]. Shadden et al. [2] later defined LCS as ridges of the finite time Lyapunov exponent (FTLE)

field. Throughout this paper we will use this definition of LCS (ridges of the FTLE field), but other definitions of LCS will be discussed in the conclusions. LCS are known to have many useful properties such as denoting barriers to transport and repelling or attracting material surfaces [2]. Additionally, they form unambiguous boundaries to well known coherent structures such as vortices [3, 4] and are relatively insensitive to small errors [5].

The FTLE field is typically computed numerically by integrating particle trajectories through the flow to approximate the flow map and then using finite differences to approximate the Jacobian of the flow map. The flow map which maps particles from their initial position at time t_0 to a final position at time $t_0 + T$, Φ , is given by

$$\Phi_{t_0}^{t_0+T}(\mathbf{x}) = \mathbf{x}(t_0) + \int_{t_0}^{t_0+T} \mathbf{v}(\mathbf{x}(t))dt \quad (1)$$

and may be computed from analytical, experimental, or numerical velocity data. If the velocity data is not analytically defined it is usually necessary to read the velocity from data files and perform interpolations in space and time. The Jacobian of Φ is used to compute the deformation tensor, Δ , which contains information about the stretching in the flow,

$$\Delta = \left(\frac{d\Phi}{d\mathbf{x}} \right)^* \left(\frac{d\Phi}{d\mathbf{x}} \right). \quad (2)$$

Finally, the largest eigenvalue of Δ is used to define the finite time Lyapunov exponent, σ ,

$$\sigma_{t_0}^T(\mathbf{x}) = \frac{1}{|T|} \ln \sqrt{\lambda_{max}(\Delta)}. \quad (3)$$

Large values of the FTLE correspond to large amounts of stretching in the flow. Additionally, one may compute the flow map either forward or backward in time (positive or negative T). Ridges in the FTLE field are then expected to correspond to either the locally most attracting or repelling lines in the flow. This is not strictly true since it is also possible for high shear regions to exhibit large FTLE values, but experience has shown that FTLE ridges are often sufficient to learn about the underlying flow structure. More advanced techniques may be used to ensure that shear structures are not selected or that the resulting LCS are exact barriers to transport (e.g. Haller, [6]).

One of the largest hurdles to more widespread use of LCS techniques is the large time required to compute the FTLE. For typical fluid flows, computing

the FTLE field requires advecting large numbers of particles through the flow. This must be done for each time step at which the FTLE field is desired. These particle advections dominate the total computational cost of any LCS algorithm.

More efficient algorithms are needed to address the large cost associated with LCS computations. Several past papers have attempted to address this problem with adaptive mesh refinement (AMR) algorithms that refine the computational mesh near the LCS [7, 8]. This is effective, but still results in computing many FTLE values away from the LCS.

Additional attempts have been made to reuse computations to compute the LCS at subsequent time steps [9]. Since the LCS are often desired at many different times and the time step may be smaller than the integration time used this leads to overlapping integrations. For example, if the LCS are desired at times $t = \{0.0, 0.1, 0.2, \dots, 10.0\}$ and an integration time of $T = 1.0$ is to be used, the standard approach would be to compute a series of flow maps $\{\Phi_{0.0}^{10.0}, \Phi_{0.1}^{10.1}, \Phi_{0.2}^{10.2}, \dots\}$ for each time the FTLE field is required. However, it is instead possible to compute the flow maps $\{\Phi_{0.0}^{0.1}, \Phi_{0.1}^{0.2}, \Phi_{0.2}^{0.3}, \dots\}$ and use the fact that $\Phi_{0.0}^{10.0} = \Phi_{9.9}^{10.0} \circ \Phi_{9.8}^{9.9} \circ \Phi_{9.7}^{9.8} \dots \Phi_{0.0}^{0.1}$. This composition of flow maps technique has the potential for large efficiency gains, but only if there is a significant overlap in the integration times and many time steps are desired. The method also comes at the cost of greatly increased memory usage to store all the necessary flow maps.

Another recent paper has reformulated the FTLE problem as an Eulerian level set problem involving the solution of a Liouville equation [10]. This allows the use of any previously developed high order accurate schemes for the resulting hyperbolic PDE's. Although there are many existing techniques for solving such hyperbolic systems, the time step used in the Eulerian method must obey a CFL condition to ensure stability while the time step in more commonly used Lagrangian techniques is typically determined by the required accuracy. This typically means that the Lagrangian techniques are faster.

We propose another alternative to any of these approaches: detect and track the ridges on the fly. A recent publication discusses a gridless ridge tracking algorithm for computing the FTLE ridges in 2D flows [11]. However, despite it's effectiveness in 2D (providing speedups of up to $80\times$), that technique is not accompanied by proofs of convergence or error bounds on the results and does not readily generalize to higher dimensions. This is because one-dimensional ridges in a 2D flow may be represented as a sim-

ple curve and need only be tracked in two directions, but in three or more dimensions, the ridges are at least two-dimensional surfaces which require a more sophisticated representation.

In this paper, we present a fast algorithm for efficiently computing LCS in three-dimensional flows. In fact, most of the algorithm may be used for $n > 0$ dimensions, but the surface triangulation used here is specific to 2D surfaces in a 3D space. All computations are performed on a predetermined orthogonal grid to simplify implementation and surface triangulation. A grid-less algorithm requires significant amounts of time to generate surface meshes, but by using a fixed grid, we are able to efficiently generate a surface triangulation via a lookup table similar to the marching cubes algorithm that is used for computing isosurfaces [12]. FTLE ridges are initially detected by computing the FTLE values on a series of lines across the domain. Local maxima along these lines occur at the FTLE ridge crossings. Once a series of points on the ridges have been detected, nearby points are tested to see if they are also on the ridge. This process is repeated to track the ridges through the entire domain. By performing computations only near the LCS surfaces, the order of the algorithm is reduced from $\mathcal{O}(1/dx^3)$ to about $\mathcal{O}(1/dx^2)$.

We present the results from several test cases, including a time dependent double gyre, Arnold-Beltrami-Childress (ABC) flow, and a swimming jellyfish. These results establish the computational order of the algorithm. We also investigate the computation cost for an FTLE field where the surface area of the LCS is known a priori and find that for a fixed grid spacing, the computational time is $C + \mathcal{O}(A_{LCS})$ where C is a constant initialization cost and A_{LCS} is the area of the LCS surfaces. Finally, the ridge tracking algorithm offers several other advantages beyond the savings in computational time. Since the LCS surfaces are directly computed, visualization of the LCS is simplified

2. A surface tracking algorithm

The computational savings seen by using a ridge tracking algorithm come from avoiding unnecessary computations away from the FTLE ridges. We separate this process into three steps: detecting initial points on the ridge surfaces, tracking the ridges through space, and triangulating the ridge points into a ridge surface. The initial ridge detection is handled by detecting where lines through the domain cross the ridge surfaces. The ridges are then iteratively grown by searching for nearby ridge points until no new ridge points




Figure 1: Initial ridge detection is handled by looking for local FTLE maxima along a few lines through the domain.

are found. Finally, the use of a gridded coordinate system throughout this process allows the use of a lookup table to efficiently generate a triangulation of the resulting LCS surfaces.

To begin, we need a ridge definition. Shadden et al. [2] offers two definitions of ridges in two dimensions. Here, we extend the concept of second derivative ridges to n dimensions:

Definition *Ridge*: A *ridge* of a C^2 function F is a co-dimension one surface S satisfying

1. The vectors $\mathbf{n} \cdot \nabla F = 0$ for all points on S where \mathbf{n} is a unit vector normal to S .
2. $\mathbf{n}^T H \mathbf{n} = \min_{\|\mathbf{u}\|=1} (\mathbf{u}^T H \mathbf{u}) < 0$ for all points on S where H is the Hessian matrix associated with F .

2.1. Initial ridge detection

If a hiker walks in a straight line, constantly monitoring his altitude, he will reach a locally maximum altitude upon crossing a ridge in the terrain. Similarly, if the FTLE is known along a line through a three-dimensional domain, local maxima along the line occur where the line crosses FTLE ridges. We restrict our computations to a fixed grid with spacing dx and initially detect the FTLE ridges by computing the FTLE values along a set of lines that cross the domain as seen in Fig. 1. The number and spacing of the lines is determined by the user and depends on the expected spatial extend of the LCS in the flow. If an LCS surface does not intersect any of the lines used and is isolated from other LCS it may be missed entirely.

Once the FTLE values along these lines are known, we look for local maxima on each line and define grid ridge points as follows:

Definition *Grid Ridge Point*: Given an orthogonal grid in \mathbb{R}^n with coordinate directions $\mathbf{e}_i : i \in \{1, \dots, n\}$; a grid point \mathbf{x}_0 is a *grid ridge point* of a function F if $F(\mathbf{x}_0) \geq F(\mathbf{x}_0 \pm \mathbf{e}_i)$ for at least one $i \in \{1, \dots, n\}$.

The local maxima on each line are grid ridge points and we begin tracking the FTLE ridges from these points. In the limit as grid spacing goes to zero, the grid ridge points converge to coordinate local maxima defined as:

Definition *Coordinate local maximum*: A point \mathbf{x}_0 is a coordinate local maximum of the function F with respect to coordinate direction \mathbf{e}_i if there is a value $\varepsilon > 0$ such that $F(\mathbf{x}_0) > F(\mathbf{x}_0 + \delta \cdot \mathbf{e}_i)$ for all $\delta < \varepsilon$.

This property is proven below in Section 3.

It is also desirable to set a threshold for the FTLE ridge values at this time. Only detecting ridges with FTLE values above some threshold ensures that only the strongest LCS are revealed. This is often done by restricting the ridges to have FTLE values above a percentage of the maximum FTLE values that are detected. Typically 50 – 70% of the maximum value is an adequate choice. Although the grid ridge points may exhibit false positives in the sense that such points may not necessarily correspond to FTLE ridges as defined by Shadden et al. [2], experience has shown that this happens infrequently and typically does not alter the topology of the resulting ridge surfaces.

2.2. Ridge tracking

The heart of this algorithm lies in tracking the ridges outward from the initially detected grid ridge points. A schematic of this process for a two dimensional example is shown in Fig. 2. The neighbors of each initially detected grid ridge point are checked to see if any meet the criteria to be grid ridge points (Fig. 2b). In 3D, points in the grid are indexed by (i, j, k) for the x , y , and z directions. For each newly detected grid ridge point, (i, j, k) , the neighboring points in $[i - 1, i + 1] \times [j - 1, j + 1] \times [k - 1, k + 1]$ are checked to see if any are grid ridge points. If new grid ridge points are found that lie above the FTLE threshold the neighbors of those points are checked. This process is repeated until no new grid ridge points are found.

It is important to note that this algorithm detects *grid ridge points* as defined above, and therefore detects *coordinate local maxima* (also defined

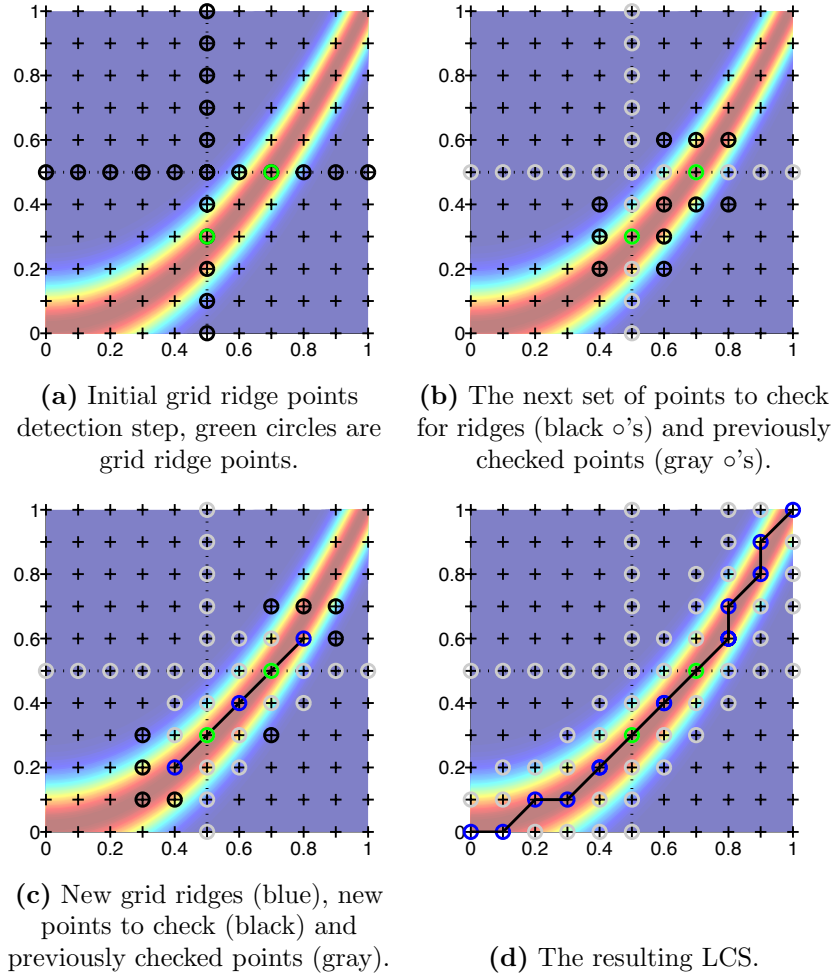


Figure 2: The ridge tracking process in 2D. The background color represents the FTLE field, +’s mark grid points. The 3D process is analogous, but a surface triangulation is used instead of line segments.

above). If the ridge height is not constant (i.e. $\nabla F \neq \mathbf{0}$ on the ridge) the grid ridge points will not converge to the actual ridges. Section 3 provides an error bound and corresponding proofs addressing this issue, but the error is typically smaller than reasonable grid spacings.

2.3. Surface triangulation

At the end of the ridge tracking process, a large list of grid ridge points is obtained. These points must be connected into a surface triangulation for visualization and further analysis. Initial attempts at developing a 3D ridge tracking algorithm revealed that grid-less techniques requiring surface meshing were both complex and expensive because of the surface meshing process. By using a gridded coordinate system, it is possible to very quickly generate a surface triangulation from a lookup table similar to the process used in the marching cubes algorithm that is popular for isosurface generation [12].

After all the grid ridge points have been found, each cubic element in the domain is given an 8-bit number corresponding to the configuration of the grid ridge points on the element's 8 vertices. The 256 possible configurations are listed in a lookup table which directly converts the 8-bit number to a triangulation of the grid ridge points in the element.

Generating the lookup table is accomplished by using reflections and rotations to reduce the 17 canonical cases shown in Fig. 3. Elements containing two or fewer grid ridge points contain no LCS surface triangles while elements containing 3 or more grid ridge points are added to the surface triangulation.

The lookup table makes the surface triangulation step in this algorithm extremely efficient, but at the cost of adding some ambiguity in the precise triangulation that should be used for certain ridge point configurations. This is a well known problem in isosurface construction and is commonly dealt with by making assumptions about the field at the subgrid scale and testing additional points within the element. Similar tests may be possible for the FLTE ridges, but we have chosen to err on the side of adding extra triangles to these ambiguous cases surfaces to avoid gaps in the LCS surfaces.

3. Algorithm properties

In this section we address the important properties of the ridge tracking algorithm by proving that well defined ridges are detected by the ridge

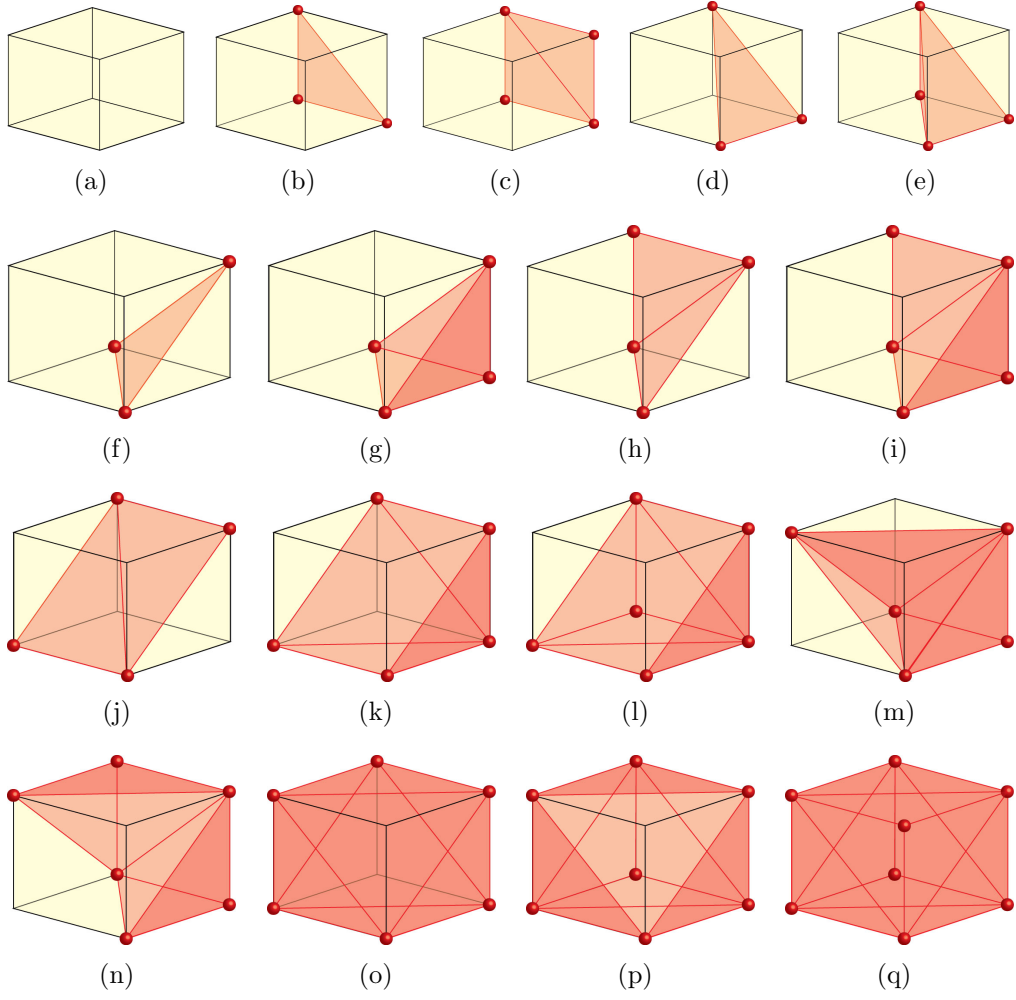


Figure 3: The canonical cases used in the lookup table for surface triangulation.

tracking algorithm and the error in the location of the ridges is typically very small. The theorems contained below rely on the ridge definition presented above in Section 2. A few important properties are summarized here for convenience:

- ∇F is parallel to the ridge.
- The eigenvectors $\{\mathbf{v}_i\}$ of the Hessian H of F form a complete, orthonormal basis for the space \mathbb{R}^n since H is symmetric.
- The eigenvector \mathbf{v}_1 associated with the minimum eigenvalue λ_1 of H is normal to the ridge.

To prove that the ridge tracking algorithm accurately detects ridges, we first show that grid ridges converge to coordinate local maxima and then show that well defined ridges always have a nearby coordinate local maximum and find a bound on the distance between a ridge and the nearest coordinate local maximum.

Theorem 3.1. *Let F be a C^2 continuous function and let L be a grid line parallel to \mathbf{e}_i . Assume \mathbf{x}^* on L is a local maximum of F in the \mathbf{e}_i direction and the second derivative of F in the \mathbf{e}_i direction is negative. Then there exists a value ϵ such that if the grid spacing $dx < \epsilon \implies \|\mathbf{x}^* - \mathbf{x}^g\| \leq dx$ for some grid point \mathbf{x}^g*

Proof. Since \mathbf{x}^* is a local maximum in the \mathbf{e}_i direction, there is a value $\varepsilon_1 > 0$ such that for all $|c_1| < \varepsilon_1$, $F|_{\mathbf{x}^* + c_1 \mathbf{e}_i} < F|_{\mathbf{x}^*}$. Since F is C^2 continuous and $\partial^2 F / \partial \mathbf{e}_i^2 < 0$, $\exists \varepsilon \in (0, \varepsilon_1)$ such that

$$\left. \frac{\partial^2 F}{\partial \mathbf{e}_i^2} \right|_{\mathbf{x}^* + c \mathbf{e}_i} < 0 \quad \forall \quad |c| < \varepsilon. \quad (4)$$

This implies that $\partial F / \partial \mathbf{e}_i$ is monotonically decreasing over the interval $c \in (-\varepsilon, \varepsilon)$.

Next, choose a grid spacing, $dx < \varepsilon = \varepsilon/2$. This ensures that at least five grid points lie in the interval $(\mathbf{x}^* - \varepsilon \mathbf{e}_i, \mathbf{x}^* + \varepsilon \mathbf{e}_i)$ and at least lie two to each side of \mathbf{x}^* .

If \mathbf{x}^* is a grid point, then the neighboring grid points have lower values of F since $dx < \varepsilon < \varepsilon_1$ so this grid point is a grid ridge point and there is a grid ridge point less than dx from \mathbf{x}^* .

If \mathbf{x}^* is not a grid point, label the nearest four grid points $\mathbf{x}_1, \mathbf{x}, \mathbf{x}_3, \mathbf{x}_4$, with \mathbf{x}^* lying between \mathbf{x} and \mathbf{x}_3 . Since $\partial F/\partial \mathbf{e}_i$ is monotonically decreasing over the interval and $\partial F/\partial \mathbf{e}_i|_{\mathbf{x}^*} = 0$, $F|_{\mathbf{x}_1} < F|_{\mathbf{x}}$ and $F|_{\mathbf{x}_3} > F|_{\mathbf{x}_4}$. If $F|_{\mathbf{x}} > F|_{\mathbf{x}_3}$, \mathbf{x} is a grid ridge point. If $F|_{\mathbf{x}} < F|_{\mathbf{x}_3}$, \mathbf{x}_3 is a grid ridge point. If $F|_{\mathbf{x}} = F|_{\mathbf{x}_3}$, both \mathbf{x} and \mathbf{x}_3 are grid ridge points. In any of these cases, since \mathbf{x}^* lies between \mathbf{x} and \mathbf{x}_3 , there is a grid ridge point less than dx from \mathbf{x}^* . \square

Thm. 3.1 proves that as grid spacing goes to zero, there are grid ridge points that converge to the coordinate local maxima of F . Therefore the ridge tracking algorithm detects coordinate local maxima. We now show that every well defined ridge has a nearby coordinate local maximum.

Theorem 3.2. *Given a C^3 continuous function F that admits a sufficiently sharp second derivative ridge, for every point \mathbf{x}_0 on the ridge, there is a nearby point, \mathbf{x}^* , in the ridge normal direction that is a coordinate local maximum. The nearest coordinate local maximum is no further from the ridge than*

$$d = \frac{2\sqrt{n-1} \|\mathbf{D}F(\mathbf{x}_0)\|}{|\lambda_1|},$$

as long as

$$\|\mathbf{v}_1 \cdot (\mathbf{D}^3 F)\| < \frac{1}{2} \frac{\lambda_1^2}{n\sqrt{n-1} \|\mathbf{D}F(\mathbf{x}_0)\|}$$

and

$$\lambda_n \leq \frac{n}{n-1} \left| \frac{\lambda_1}{n} + \frac{2\sqrt{n-1}}{|\lambda_1|} \|\mathbf{D}F(\mathbf{x}_0)\| \|\mathbf{v}_1 \cdot \mathbf{D}F(\xi)\| \right|,$$

where \mathbf{D} denotes the gradient operator, λ_1 (respectively λ_n) is the minimum (respectively maximum) eigenvalue of the Hessian, $\mathbf{D}^2 F(\mathbf{x}_0)$, \mathbf{v}_1 and \mathbf{v}_n are the eigenvectors associated with λ_1 and λ_n , and n is the dimension of the space.

Although these criteria may seem restrictive, they are typically easily satisfied by well defined ridges as we will see in examples below. λ_1 is typically very large in magnitude (sometimes up to 10^{12}) while $\|\mathbf{D}F\|$ and $|\lambda_n|$ are typically $\mathcal{O}(1)$. If λ_n is negative the third condition is trivially satisfied.

Proof. By assumption F is C^3 continuous and admits a second derivative ridge through point \mathbf{x}_0 . Denote the eigenvalues and normalized eigenvectors

of the Hessian $\mathbf{D}^2 F(\mathbf{x}_0)$ as $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ and $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$. By definition, \mathbf{v}_1 is normal to the ridge and $\mathbf{D}F(\mathbf{x}_0)$ is parallel to the ridge. We also know that $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ forms an orthonormal basis for the space since the Hessian is a real symmetric matrix. Also, let the coordinate system for the space be defined by the orthonormal basis vectors $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$.

We first choose the coordinate direction that is closest to the ridge normal direction. That is, we choose \mathbf{e}_j from the set $\{\mathbf{e}_i\}$ such that $|\mathbf{e}_j \cdot \mathbf{v}_1| \geq |\mathbf{e}_i \cdot \mathbf{v}_1| \forall i \in \{1, \dots, n\}$. By Thm. Appendix A.1 we know that

$$|\mathbf{e}_j \cdot \mathbf{v}_1| \geq 1/\sqrt{n}. \quad (5)$$

The gradient of F near \mathbf{x}_0 may be written as

$$\mathbf{D}F(\mathbf{x}) = \mathbf{D}F(\mathbf{x}_0) + ((\mathbf{x} - \mathbf{x}_0) \cdot \mathbf{D}) \mathbf{D}F(\mathbf{x}_0) + \frac{1}{2!} ((\mathbf{x} - \mathbf{x}_0) \cdot \mathbf{D})^2 \mathbf{D}F(\xi)$$

for some ξ that is a linear combination of \mathbf{x} and \mathbf{x}_0 . Taking the dot product of $\mathbf{D}F(\mathbf{x})$ with \mathbf{e}_j and considering only points on a line normal to the ridge such that $\mathbf{x} = \mathbf{x}_0 + d\mathbf{v}_1$ gives

$$\mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}) = \mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}_0) + d\mathbf{e}_j \cdot \{(\mathbf{v}_1 \cdot \mathbf{D}) \mathbf{D}F(\mathbf{x}_0)\} + \frac{1}{2!} d^2 \mathbf{e}_j \cdot \{(\mathbf{v}_1 \cdot \mathbf{D})^2 \mathbf{D}F(\xi)\}.$$

Since \mathbf{v}_1 is an eigenvalue of $\mathbf{D}^2 F(\mathbf{x}_0)$, this can be rewritten as

$$\mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}) = \mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}_0) + d\lambda_1 \mathbf{e}_j \cdot \mathbf{v}_1 + \frac{1}{2!} d^2 \mathbf{e}_j \cdot (\mathbf{v}_1 \cdot \{\mathbf{D}^3 F(\xi)\}) \quad (6)$$

We can establish an upper bound on $\mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}_0)$ by noting that $\mathbf{v}_1 \cdot \mathbf{D}F(\mathbf{x}_0) = 0$ and expressing \mathbf{e}_j and $\mathbf{D}F(\mathbf{x}_0)$ in terms of $\{\mathbf{v}_i\}$:

$$\begin{aligned} |\mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}_0)| &= \left(\sum_{i=1}^n (\mathbf{v}_i \cdot \mathbf{e}_j) \mathbf{v}_i \right) \cdot \left(\sum_{i=2}^n (\mathbf{v}_i \cdot \mathbf{D}F(\mathbf{x}_0)) \mathbf{v}_i \right) \\ &= \left(\sum_{i=2}^n (\mathbf{v}_i \cdot \mathbf{e}_j) \mathbf{v}_i \right) \cdot \left(\sum_{i=2}^n (\mathbf{v}_i \cdot \mathbf{D}F(\mathbf{x}_0)) \mathbf{v}_i \right) \\ &\leq \left\| \sum_{i=2}^n (\mathbf{v}_i \cdot \mathbf{e}_j) \mathbf{v}_i \right\| \left\| \sum_{i=2}^n (\mathbf{v}_i \cdot \mathbf{D}F(\mathbf{x}_0)) \mathbf{v}_i \right\| \\ &\leq \sqrt{\sum_{i=2}^n (\mathbf{v}_i \cdot \mathbf{e}_j)^2} \sqrt{\sum_{i=2}^n (\mathbf{v}_i \cdot \mathbf{D}F(\mathbf{x}_0))^2} \\ &\leq \sqrt{1 - \frac{1}{n}} \|\mathbf{D}F(\mathbf{x}_0)\|. \end{aligned} \quad (7)$$

The right hand side of Eq. 6 is a quadratic function in d with bounds on the coefficients

$$\begin{aligned} 0 \leq |c| \leq c_1 &= \sqrt{1 - 1/n} \|\mathbf{D}F(\mathbf{x}_0)\|, \\ 0 < b_0 = \frac{|\lambda_1|}{\sqrt{n}} \leq |b| \leq b_1 &= |\lambda_1|, \\ 0 \leq |a| \leq \frac{\lambda_1^2}{4n\sqrt{n-1} \|\mathbf{D}F(\mathbf{x}_0)\|} &< \frac{b_0^2}{4c_1}. \end{aligned}$$

Therefore, by Thm. Appendix A.2, this function has a root at d^* with

$$|d^*| < \frac{2\sqrt{n-1} \|\mathbf{D}F(\mathbf{x}_0)\|}{\lambda_1}.$$

where $\mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}) = 0$.

We have now shown that the value of $\mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x})$ is zero somewhere on the line normal to the ridge at \mathbf{x}_0 at a distance of less than

$$2\sqrt{(1 - 1/n)} \|\mathbf{D}F(\mathbf{x}_0)\| / \lambda_1.$$

Call the zero point \mathbf{x}^* . To complete the proof, we will show that $\mathbf{e}_j^T(\mathbf{D}^2F)\mathbf{e}_j$ must be negative at this point so it is a coordinate local maximum. The first order Taylor series for the Hessian near \mathbf{x}_0 is

$$\mathbf{D}^2F(\mathbf{x}) = \mathbf{D}^2F(\mathbf{x}_0) + ((\mathbf{x} - \mathbf{x}_0) \cdot \mathbf{D})\mathbf{D}^2F(\xi).$$

At \mathbf{x}^* the second derivative of F in the \mathbf{e}_j direction is given by

$$\mathbf{e}_j^T(\mathbf{D}^2F(\mathbf{x}^*))\mathbf{e}_j = \mathbf{e}_j^T(\mathbf{D}^2F(\mathbf{x}_0))\mathbf{e}_j + d^*\mathbf{e}_j^T(\mathbf{v}_1 \cdot \mathbf{D}^3F(\xi))\mathbf{e}_j$$

where $d^* < 2\sqrt{n-1} \|\mathbf{D}F(\mathbf{x}_0)\| / \lambda_1$ is the distance from the ridge of the point where $\mathbf{e}_j \cdot \mathbf{D}F(\mathbf{x}) = 0$. Then

$$\begin{aligned} \mathbf{e}_j^T\mathbf{D}^2F(\mathbf{x}) &< \frac{\lambda_1}{n} + \left(1 - \frac{1}{n}\right) \lambda_n + \frac{2\sqrt{n-1} \|\mathbf{D}F(\mathbf{x}_0)\|}{|\lambda_1|} \|\mathbf{v}_1 \cdot \mathbf{D}^3F(\xi)\| \\ &< 0 \end{aligned}$$

where $\|\mathbf{v}_1 \cdot \mathbf{D}^3F(\xi)\|$ is understood as the induced norm of the matrix that results from the tensor dot product $\mathbf{v}_1 \cdot \mathbf{D}^3F(\xi)$. Thus, at \mathbf{x}^* the first derivative of F is zero in the \mathbf{e}_j direction and the second derivative is negative. \mathbf{x}^* is a coordinate local maximum. □

As long the criteria of Thm. 3.1 are satisfied along a ridge, there is a nearby coordinate local maximum that will be detected by the ridge tracking algorithm. The error of the location of the detected ridges is bounded by

$$d = \frac{2\sqrt{n-1} \|\mathbf{D}F(\mathbf{x}_0)\|}{|\lambda_1|}.$$

Since $\|\mathbf{D}F\|$ is typically $\mathcal{O}(1)$ (if the ridge does not rise and fall too quickly), the error is typically $\mathcal{O}(1/|\lambda_1|)$.

3.1. Ridge examples

We now verify the properties of the ridge tracking algorithm on two two-dimensional examples. An analytically defined surface that admits a ridge and the FTLE field for a time dependent double gyre flow. We begin by examining the surface

$$F(x, y) = \frac{e^{-0.1(x-y)^2}}{\ln((x+y)^2 + 15)}. \quad (8)$$

A contour plot of this function is shown in Fig. 4 and the function admits a ridge along the line $y = x$. This represents a worst case scenario in terms of ridge orientation since the ridge is at a 45° angle to both coordinate directions. The ridge is also much thicker than is typically seen in FTLE ridges. Since this ridge is defined by an analytical function, it is possible to easily compute all the necessary criteria for Thm. 3.2. In Fig. 4, the actual ridge has been drawn as a solid black line. The error bound, d is drawn in the figure as a pair of dashed black lines and the coordinate local maxima have been drawn as red (for $\partial F/\partial x = 0$) and blue (for $\partial F/\partial y = 0$) lines. The other criteria of Thm. 3.2 are all satisfied despite λ_1 being much larger than is usually seen in FTLE ridges. We find $\lambda_1 \in [-0.1477, -0.0843]$,




Figure 4: The surface defined by Eq. 8. The solid black line denotes the ridge while the dashed black curves show the error bound and the red and blue curves show the coordinate local maxima.




Figure 5: (a) The forward FTLE field and (b) the main FTLE ridge (with FTLE value > 0.25) for the time dependent double gyre with $A = 0.1$, $\epsilon = 0.25$, $\omega = 10$, $t = 0$, and $T = 15$.

$\lambda_2 \in [-0.0363, 0.0071]$, and $\|\mathbf{D}F\| \leq 0.0356$ on the ridge in the domain $(x, y) \in [-5, 5]^2$.

As a second example to apply the above theorem, we consider a ridge in the FTLE field of a time dependent double gyre. The velocity field consists of two counter rotating gyres with a periodic perturbation that enables transport between the two gyres. The flow is given by the stream function

$$\psi(x, y, t) = A \sin(\pi f(x, t)) \sin(\pi y) \quad (9)$$

on the domain $[0, 2] \times [0, 1]$ where

$$\begin{aligned} f(x, t) &= a(t)x^2 + b(t), \\ a(t) &= \epsilon \sin(\omega t), \\ b(t) &= 1 - 2\epsilon \sin(\omega t). \end{aligned} \quad (10)$$

The velocity is given by

$$u = -\frac{\partial \psi}{\partial y}, \quad v = \frac{\partial \psi}{\partial x}. \quad (11)$$

We use parameters $A = 0.1$, $\epsilon = 0.25$, and $\omega = 10$ for this example and set the integration time for computing the FTLE at $T = 15$. We will only consider the forward time FTLE field at time $t = 0$ which is shown in Fig. 5.

The main FTLE ridge in this double gyre flow (shown in Fig. 5b) was computed with very high precision by iteratively estimating the ridge position and tangent direction and then adjusting the position in the normal direction.




Figure 6: (a) The FTLE values and gradient along the ridge shown in Fig. 5b and (b) bound on the ridge position error as given in Thm. 3.2. The distance along the ridge (x -axis) is computed as the distance along the ridge from the ridge origin at $(1.110, 0)$. Since $\|\nabla F\| = \mathcal{O}(1)$ and $|\lambda_1| \gg 1$, the error bound is very small, typically less than 10^{-6} .

This is necessary to accurately compute the FTLE values on the ridge as well as the gradient, and Hessian of the FTLE field to bound the error in the ridge locations. Note that the ridge seen in this FTLE field is much sharper than the analytical ridge investigated in the previous example.

Fig. 6 shows the FTLE values and the norm of the gradient of the FTLE field along the ridge as well as the bound on the ridge location error which is less than 10^{-6} for the majority of the ridge and never rises above 10^{-2} . The norm of the gradient is bounded by $\|\nabla F\| < 1.5$. Additionally the eigenvalues of the Hessian fall in the range $-1.8 \times 10^{12} < \lambda_1 < -367$ and $-6.1 < \lambda_2 < 37.3$ with averages of $\bar{\lambda}_1 = -8.1 \times 10^{10}$ and $\bar{\lambda}_2 = -0.033$ and medians $\text{med}(\lambda_1) = -5.71 \times 10^7$ and $\text{med}(\lambda_2) = -0.30$. The very large magnitude of λ_1 means that the grid based ridge tracking algorithm will be extremely accurate for this example. For reference, grid spacings of 10^{-2} or 10^{-3} are typically used when performing LCS computations for this problem.

4. Algorithm performance

In this section we discuss the performance of the ridge tracking algorithm by examining two analytically defined examples: a time dependent double gyre and Arnold-Beltrami-Childress flow. We verify that the expected LCS surfaces are extracted and then focus on establishing the computational order of the surface tracking algorithm and compare this to the standard FTLE algorithm that computes the FTLE field everywhere in the domain.




Figure 7: Surface tracking results for the double gyre flow. The 3D LCS are shown at in (a) with forward LCS colored blue and backward LCS colored red. The forward FTLE field is shown in (b) with the ridge tracking results overlaid as the black curve that precisely lines up with the FTLE ridge.

4.1. Time dependent double gyre

The first example presented is the time dependent double gyre used above in section 3.1. We extend this flow to three dimensions by simply setting the velocity in the z direction to $w = 0$. Since there is no z dependence and no velocity in the z direction, the LCS will be independent of z as well. We use the parameters $A = 0.1$, $\epsilon = 0.1$, and $\omega = 2\pi/10$. We set the integration time to be $T = \pm 15$ and compute both the forward and backward LCS. A threshold of 80% of the maximum FTLE value is used to determine the LCS.

The LCS in this system have been well studied in the past and our results agree with previous publications [2, 11]. The full 3D LCS surfaces are shown in Fig. 7. This figure also shows the backward FTLE field overlaid with the results of the 3D ridge tracking algorithm. The LCS extracted by the ridge tracking algorithm lie exactly on top of the ridge in the FTLE field.

The computational timing results are summarized below in Section 4.4

4.2. Arnold-Beltrami-Childress flow

Arnold-Beltrami-Childress (ABC) flow is a three-dimensional, 2π -periodic flow that has been previously studied with LCS techniques [13]. The flow is given by

$$\begin{aligned} u &= A \sin(z) + C \cos(y), \\ v &= B \sin(x) + A \cos(z), \\ w &= C \sin(y) + B \cos(x), \end{aligned} \tag{12}$$

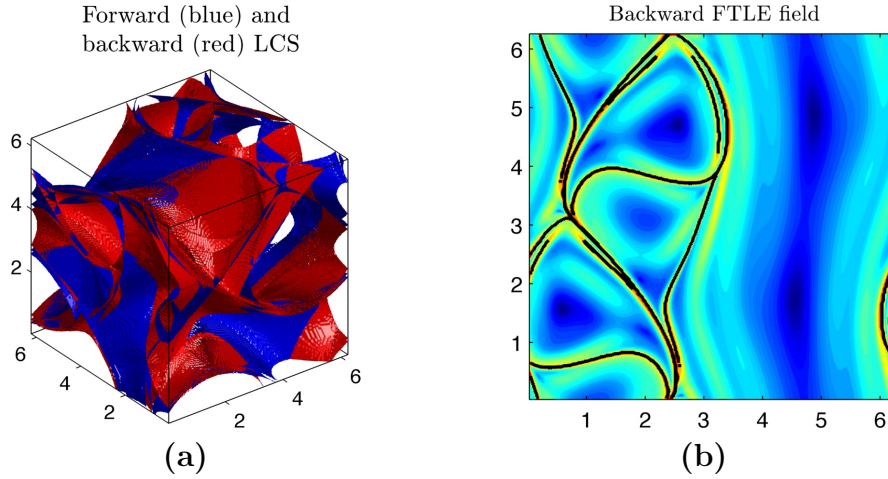


Figure 8: Ridge tracking results for ABC flow. The 3D LCS are shown in (a) with forward LCS colored blue and backward LCS colored red. (b) shows a single plane at $z = \pi$ to display the FTLE field (colored) and corresponding LCS (black curves) as computed with the ridge tracking algorithm.

where $A = 1$, $B = \sqrt{2/3}$, and $C = \sqrt{1/3}$. We use an integration time of $T = 10$ for FTLE computations and use a threshold value of 70% of the maximum FTLE value.

The LCS are shown in Fig. 8 and agree with previously published results [13, 14]. Fig. 8 also shows single slice of the FTLE field well as the FTLE field and corresponding LCS at a height of $z = \pi$. The complex LCS present in the ABC flow are a product of the non-trivial invariant manifolds of this flow. They clearly divide the flow into different regions that appear as tube-like structures through the flow domain. These tubes are dynamically distinct from one another and particles travel within and along the tubes without escaping to other regions of the space.

4.3. Swimming jellyfish

As a final example, we compute the LCS created by a jetting type jellyfish. The jellyfish body motion was extracted from a digital video of a swimming individual jellyfish (species *Sarsia tubulosa*) and use as input to an arbitrary Lagrangian-Eulerian CFD code that computes the resulting flow field and jellyfish acceleration. Full details of this procedure can be found in Sahin and Mohseni [15] and Sahin et al. [16].

The resulting axisymmetric velocity field is returned on a moving, non-


uniform quadrilateral mesh in (r, z) -coordinates. During LCS computations, the (x, y, z) coordinates are converted to (r, z) coordinates to compute the velocity and the velocity is converted back to Cartesian coordinates for particle advections. The mesh type creates significant complications for velocity interpolation during particle advection since even locating the mesh element that contains a given point is non-trivial. To address this issue efficiently, an alternating digital tree (ADT) is used to search the domain for the element that contains a given drifter particle [17]. The ADT recursively divides the space in half so that at each node in the tree, only one branch must be searched. Since the nodes of the mesh elements are listed in counter clockwise order and the elements are convex, a point \mathbf{p} is inside (including the boundary) an element if and only if

$$\hat{\mathbf{z}} \cdot [(\mathbf{v}_i - \mathbf{p}) \times (\mathbf{v}_{(i \bmod 4)+1} - \mathbf{v}_i)] \geq 0 \quad \forall i \in \{1, 2, 3, 4\}$$


for vertices $\{\mathbf{v}_i\}$.

Once the element containing a particle is found, the velocity must be interpolated onto the drifter particle. Since the elements are generally none rectangular, simple linear interpolation is not possible. Instead, perspective projection is used to map the quadrilateral element and the point of interest onto the unit square. Bilinear interpolation is then used to approximate the velocity of the particle.

The jellyfish chosen for this investigation has a swimming period of 1 s between contractions and an integration time of 0.5 s was used to compute the LCS. Fig. 9 shows the LCS as computed with the ridge tracking algorithm as well as the backward FTLE field computed with the standard FTLE algorithm. The backward LCS (Fig. 9a) clearly show a strong vortex being ejected as the jellyfish's bell contraction comes to an end. Additionally, the forward LCS outline fluid ahead of the vortex that will soon be entrained into the vortex ring as well as a region of fluid near the jellyfish bell that will be drawn into the bell during the relaxation phase of jellyfish swimming. These results are in excellent agreement with previously published LCS for this jellyfish [18]. Furthermore, this example clearly demonstrates the visualization advantages offered by the ridge tracking algorithm. Since the actual LCS surface are computed it is much simpler to visualize both the forward and backward LCS simultaneously.



(a)



(b)

Figure 9: (a) The results of the ridge tracking algorithm and (b) the backward FTLE field for the swimming jellyfish. A strong vortex is being ejected near the end of the jellyfish's bell contraction. A cutaway view is shown so that the full LCS structure is visible.

4.4. Timing results

We expect the computational time of the ridge tracking algorithm to be $\mathcal{O}(1/dx^2)$ for a grid of spacing dx since computations are performed only near the 2D ridge surfaces. To establish the computational order, we have computed the full FTLE field using a standard FTLE algorithm and also computed the LCS with the ridge tracking algorithm presented above for the double gyre, ABC, and jellyfish flows. Computations were performed on at a variety of grid resolutions as well as on a single core and with a parallel code running on 16 or 48 cores. A least squares best fit was performed on a log-log scale for each case, assuming a fit of

$$t_f = C/(dx^\alpha). \quad (13)$$

The resulting data points and curve fits are shown in Fig. 10 and show that the standard algorithm is $\approx \mathcal{O}(1/dx^{3.0})$ while the ridge tracking algorithm scales approximately as $\mathcal{O}(1/dx^{2.1})$. This performance is maintained for the parallel version of the code.

It is worth noting that in the jellyfish example, for low resolutions the time required to read the velocity data files and build the ADTs used for search at each time step is a significant part of the total computational time. Since the velocity read in time does not change with grid resolution, this has the effect of giving artificially low exponents for the algorithm order (both for the standard algorithm and the ridge tracking algorithm). To compensate for this effect and more accurately estimate the asymptotic order, we estimated the velocity read in and ADT creation time and subtracted this value from the total run time before computing the computational order for the jellyfish example. These modified times are reported in Fig. 10d. The resulting values for α closely match the values that would result from using only the last few data points and represent a closer approximation of the asymptotic values of α .

It is also expected that the computational time should be directly proportional to the surface area of the LCS in the domain. We test this by generating an artificial FTLE field that has ridges along pre-selected planes. The planes are defined by $z = 0.1(x + y) + h_i$ and the artificial FTLE field is given by

$$\sigma(x, y, z) = \sum_i \exp[-2000(z - 0.1(x + y) - h_i)^2]. \quad (14)$$

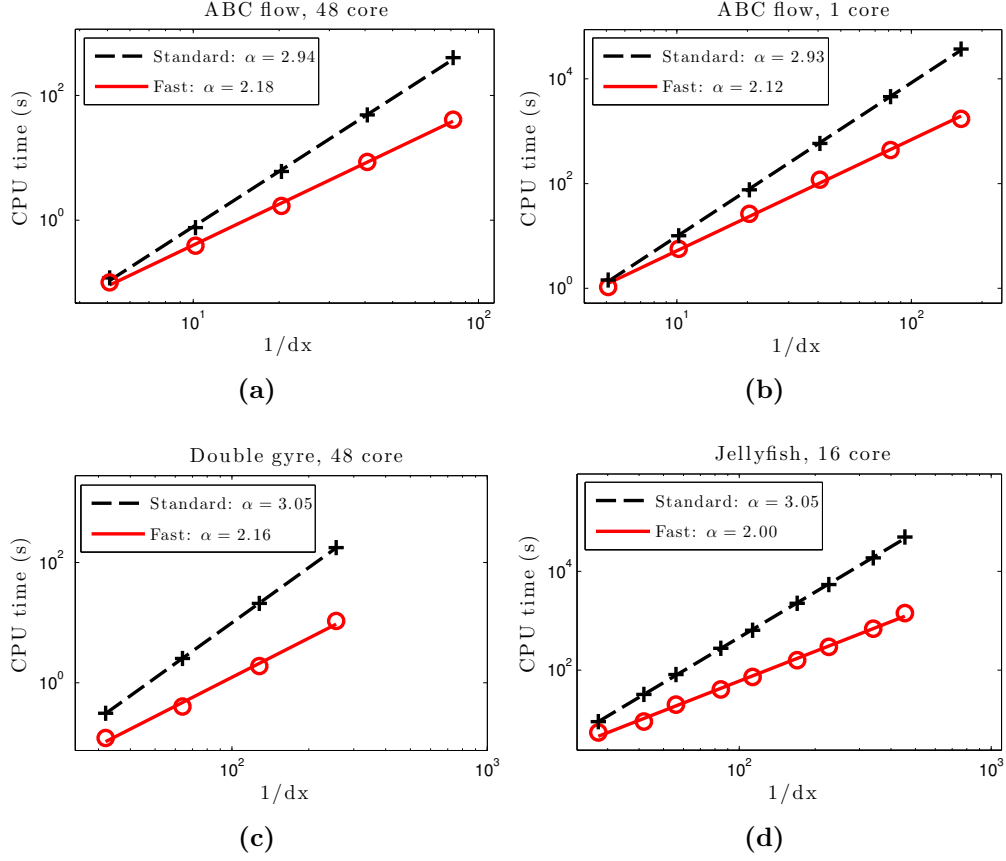


Figure 10: Timing results for the standard FTLE algorithm and the ridge tracking algorithm. α is the scaling exponent of the algorithm ($\text{CPU time} \approx \mathcal{O}(1/dx^\alpha)$) and appears as the slope of the lines in these log-log plots. The standard algorithm scales as $\mathcal{O}(1/dx^{3.0})$ and the ridge tracking algorithm scales as $\mathcal{O}(1/dx^{2.1})$.

CPU time (s)		Surface area					
		1.01	2.02	3.03	4.04	5.05	6.06
Grid Spacing	1/64	0.533	0.689	0.844	0.987	1.162	1.333
	1/128	2.713	3.306	3.913	4.625	5.385	6.121
	1/256	19.54	22.12	25.10	27.68	32.79	34.80
	1/512	192.8	204.5	215.7	235.7	244.7	259.9

Table 1: CPU time (in seconds) to compute the artificial FTLE ridges of Eq. 14 for various surface areas and grid spacings.

Each plane that defines an FTLE ridge is tilted slightly so that it is out of line with the grid to make detecting and tracking the ridge slightly more realistic. We use the domain $[0, 1]^3$ where each ridge has an area of $A_{LCS} = \sqrt{1.02}$ for $h_i \in (0, 0.8)$ and test six different cases corresponding to 1 – 6 ridges in the domain. Particles are advected using the double gyre velocity field listed above to account for the particle advection time, but instead of returning the true FTLE field for the double gyre flow, the artificial field of Eq. 14 with pre-selected ridges is returned. The results are listed in Table 1 which shows the computational time according to surface area for four different values of dx .

All four values of dx show a linear relationship between LCS surface area and CPU time and least squares fits result in the following regression coefficients for the fit $t_{CPU} = C_1 + C_2 A_{LCS}$ where A_{LCS} is the surface area of the LCS and t_{CPU} is the required CPU time:

dx	C_1	C_2
1/64	0.368	0.157
1/128	1.94	0.679
1/256	15.9	3.14
1/512	177.9	13.5

The relatively large values of C_1 for all these cases means that there is some initial cost regardless of the amount of LCS surface area. This is due to the cost associated with initializing data structures and performing the initial ridge detection step as described in Section 2.1. Additionally, C_1 appears to scale roughly as $\mathcal{O}(1/dx^3)$. Since the current implementation of the ridge tracking code uses and allocates full 3D arrays rather than using sparse data structures, it is reasonable to expect this relationship. On the other hand,

C_2 scales roughly as $\mathcal{O}(1/dx^2)$. This accounts for the majority ridge tracking part of the algorithm.

The $\mathcal{O}(1/dx^{3.0})$ scaling of C_1 may explain why the overall computational order of the algorithm is $\mathcal{O}(1/dx^{2.1})$ rather than $\mathcal{O}(1/dx^2)$. As resolution is increased beyond current capabilities, the initialization cost of may be expected to completely overwhelm the ridge tracking cost due to this difference in order. The implementation of sparse data structures would likely help solve this problem.

5. Conclusions

As the problems being analyzed with LCS techniques become increasingly complex, the corresponding computations become increasingly expensive. The ridge tracking algorithm presented in this paper has been shown to reduce the order of LCS computations from $\mathcal{O}(1/dx^{3.0})$ to about $\mathcal{O}(1/dx^{2.1})$ for three-dimensional flows. This reduction in order allows potentially tremendous savings in computational time as the required LCS resolution is increased.

The effectiveness and algorithm properties have been demonstrated by several examples, including the analytically defined double gyre and ABC flow and the swimming jellyfish that is defined by velocity stored in data files. On single processor as well as multicore machines, the ridge tracking algorithm shows the expected change in computational order and provides large speed ups for all tested cases.

We have also proved that although the ridge tracking algorithm detects *coordinate local maxima* rather than the actual ridges, for well defined ridges the associated error is small. The distance between a ridge and the surfaces detected by this algorithm is $\mathcal{O}(\|\mathbf{D}F\| / |\lambda_1|)$ where $\mathbf{D}F$ is the gradient along the ridge and $|\lambda_1|$ is the second derivative normal to the ridge (or, equivalently the smallest eigenvalue of the Hessian of F). In typical examples this error is at smaller than the grid spacing and for well defined ridges it may be several orders of magnitude smaller than the grid spacing.

The general framework of this algorithm could easily be adapted to other LCS definitions or techniques. For example, the finite size Lyapunov exponent (FSLE) could easily be used in lieu of the FTLE. A 2D version of this algorithm (or the grid-less algorithm of Lipinski and Mohseni [11]) could be relatively easily modified to use the variational formulation for LCS which was recently proposed by Haller [6, 19].

Future work will focus on implementing sparse data structures to further reduce the computational cost of large LCS calculations and on implementing some additional LCS techniques such as those mentioned in the previous paragraph. We expect that the use of sparse data structures may further reduce the order of computation from $\mathcal{O}(1/dx^{2.1})$ nearer to $\mathcal{O}(1/dx^{2.0})$.

References

- [1] G. Haller, G. Yuan, Lagrangian coherent structures and mixing in two dimensional turbulence, *Physica D* 147 (2000) 352–370.
- [2] S. Shadden, F. Lekien, J. Marsden, Definition and properties of Lagrangian coherent structures from finite time Lyapunov exponents in two-dimensional aperiodic flows, *Physica D* 212 (2005) 271–304.
- [3] S. Shadden, J. Dabiri, J. Marsden, Lagrangian analysis of fluid transport in empirical vortex ring flows, *Physics of Fluids* 18 (2006) 7105–7120.
- [4] B. Cardwell, K. Mohseni, Vortex shedding over two-dimensional airfoil: Where do the particles come from?, *AIAA J.* 46 (2008) 545–547.
- [5] G. Haller, Lagrangian coherent structures from approximate velocity data, *Phys. Fluids A* 14 (2002) 1851–1861.
- [6] G. Haller, A variational theory of hyperbolic Lagrangian coherent structures, *Physica D* 240 (2011) 574–598.
- [7] C. Garth, F. Gerhardt, X. Tricoche, H. Hagen, Efficient computation and visualization of coherent structures in fluid flow applications, *IEEE Transactions on Visualization and Computer Graphics* 13 (2007) 1464.
- [8] F. Sadlo, R. Peikert, Efficient visualization of Lagrangian coherent structures by filtered AMR ridge extraction, *IEEE Transactions on Visualization and Computer Graphics* 13 (2007) 1456.
- [9] S. Brunton, C. Rowley, Fast computation of finite-time lyapunov exponent fields for unsteady flows, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 20 (2010) 017503.
- [10] S. Leung, An Eulerian approach for computing the finite time Lyapunov exponent, *J. Comp. Physics* 230 (2011) 3500 – 3524.

- [11] D. Lipinski, K. Mohseni, A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures, *Chaos* 20 (2010) 017504 (9pp.). Doi:10.1063/1.3270049.
- [12] W. Lorensen, H. Cline, Marching cubes: A high resolution 3D surface construction algorithm, *SIGGRAPH Comput. Graph.* 21 (1987) 163–169.
- [13] G. Haller, Distinguished material surfaces and coherent structures in three-dimensional fluid flows, *Physica D* 149 (2001) 248–277.
- [14] G. Froyland, K. Padberg, Almost-invariant sets and invariant manifolds – connecting probabilistic and geometric descriptions of coherent structures in flows, *Physica D* 238 (2009) 1507–1523.
- [15] M. Sahin, K. Mohseni, An arbitrary Lagrangian-Eulerian formulation for the numerical simulation of flow patterns generated by the hydromedusa *Aequorea victoria*, *J. Comp. Phys.* 228 (2009) 4588–4605.
- [16] M. Sahin, K. Mohseni, S. Colins, The numerical comparison of flow patterns and propulsive performances for the hydromedusae *Sarsia tubulosa* and *Aequorea victoria*, *J. Exp. Biology* 212 (2009) 2656–2667.
- [17] J. Bonet, J. Peraire, An alternating digital tree (ADT) algorithm for 3D geometric searching and intersection problems, *International Journal for Numerical Methods in Engineering* 31 (1991) 1–17.
- [18] D. Lipinski, K. Mohseni, A numerical investigation of flow structures and fluid transport with applications to feeding for the hydromedusae *Aequorea victoria* and *Sarsia tubulosa*, *J. Exp. Biology* 212 (2009) 2436–2447.
- [19] M. Farazmand, G. Haller, Computing Lagrangian coherent structures from variational LCS theory, 2011. Preprint. Online at <http://georgehaller.com/reprints/compvariLCS.pdf>.

Appendix A. Additional proofs

This appendix contains additional theorems and proofs referenced by the proofs of the theorems in the text.

Theorem Appendix A.1. *Given an orthonormal basis $\{\mathbf{e}_i\}$ for \mathbb{R}^n and an arbitrary unit vector \mathbf{v} there is a basis vector \mathbf{e}_j that maximizes $|\mathbf{v} \cdot \mathbf{e}_i|$ and $|\mathbf{v} \cdot \mathbf{e}_j| \geq 1/\sqrt{n}$.*

Proof. Existence of the maximum is implied by the extreme value theorem since the dot product is a continuous function and the set $\{\mathbf{e}_i\}$ is finite and therefore closed and bounded. The other part of the proof is as follows: Assume $|\mathbf{v} \cdot \mathbf{e}_j| < 1/\sqrt{n}$. Then

$$\begin{aligned}
1 = \|\mathbf{v}\| &= \left\| \sum_{i=1}^n (\mathbf{e}_i \cdot \mathbf{v}) \mathbf{e}_i \right\| \\
&= \sqrt{\left(\sum_{i=1}^n (\mathbf{e}_i \cdot \mathbf{v}) \mathbf{e}_i \right) \cdot \left(\sum_{i=1}^n (\mathbf{e}_i \cdot \mathbf{v}) \mathbf{e}_i \right)} \\
&= \sqrt{\sum_{i=1}^n (\mathbf{e}_i \cdot \mathbf{v})^2} \\
&\leq \sqrt{\sum_{i=1}^n (\mathbf{e}_j \cdot \mathbf{v})^2} = \sqrt{n(\mathbf{e}_j \cdot \mathbf{v})^2} \\
&< \sqrt{n(1/\sqrt{n})^2} \\
&< 1
\end{aligned}$$

which is a contradiction so $|\mathbf{v} \cdot \mathbf{e}_j| \geq 1/\sqrt{n}$. □

Theorem Appendix A.2. *Given a quadratic equation $y = ax^2 + bx + c$ with bounds on the real valued constants*

$$\begin{aligned}
|c| &\leq c_1 \\
0 < b_0 &\leq |b| \leq b_1 \\
|a| &\leq \frac{b_0^2}{4c_1}
\end{aligned}$$

there is a real root x^ such that $|x^*| < \frac{2c_1}{b_0}$.*

Proof. If $a = 0$, there is a single root at $x^* = -c/b$ and $|x^*| = |-c/b| < 2c_1/b_0$.

If $a \neq 0$, the roots are given by the alternate form of the quadratic formula

$$x = \frac{2c}{-b \pm \sqrt{b^2 - 4ac}}.$$

The discriminant, $\Delta = b^2 - 4ac > b_0^2 - 4\frac{b_0^2}{4c_1}c_1 = 0$ is positive so the quadratic has two real roots. Denote these roots

$$x_1 = \frac{2c}{-b + \sqrt{b^2 - 4ac}}$$

$$x_2 = \frac{2c}{-b - \sqrt{b^2 - 4ac}}.$$

If $b \leq 0$ let $x^* = x_1$, otherwise let $x^* = x_2$. Then

$$|x^*| = \frac{2|c|}{|b| - \sqrt{b^2 - 4ac}}$$

$$< \frac{2c_1}{b_0}$$

□